



## Open Research Online

### Citation

Tessem, Bjørnar; Nyre, Lars; Mesquita, Michel and Mulholland, Paul (2022). Deep Learning to Encourage Citizen Involvement in Local Journalism. In: Manninen, V.J.E.; Niemi, M.K. and Ridge-Newman, A. eds. Futures of Journalism. Palgrave Macmillan, Cham, pp. 211–226.

### URL

<https://oro.open.ac.uk/83754/>

### License

None Specified

### Policy

This document has been downloaded from Open Research Online, The Open University's repository of research publications. This version is being made available in accordance with Open Research Online policies available from [Open Research Online \(ORO\) Policies](#)

### Versions

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding

# Deep learning to encourage citizen involvement in local journalism

Bjørnar Tessem

Dept. of Information Science and Media Studies,  
University of Bergen, Postbox7802,  
N-5020 Bergen, Norway,  
bjornar.tessem@uib.no

Lars Nyre

Dept. of Information Science and Media Studies,  
University of Bergen, Postbox7802,  
N-5020 Bergen, Norway,  
Lars.Nyre@uib.no

Michel dos Santos Mesquita

M2Lab Centre for Statistical and Data Science Research  
Hamnegata 6, 5411 Stord, Norway  
michel@m2lab.org

Paul Mulholland

Knowledge Media Institute  
The Open University  
Walton Hall, Milton Keynes, UK.  
paul.mulholland@open.ac.uk

**Abstract (155 words)**

We discuss the potential of a mobile app for news tips to local newspapers to be augmented with artificial intelligence. It can be designed to encourage deliberative, consensus-oriented contributions from citizens. We presume that such an app will generate news stories from multi-modal data in the form of photos, videos, text elements and information about the location and identity of the contributor. Three scenarios are presented to show how image recognition, natural language processing, narrative construction and other AI technologies can be applied. The scenarios address three interrelated challenges for local journalism. First, news tips from readers are often of low technical quality; containing little information and poor photos. Second, peer-to-peer dialogue about local news takes place in social media instead of in the newspaper. Third, readers lack news literacy and are prone to confrontational debates and trolling. We show how advances in deep learning technology makes it possible to propose solutions to these problems.

**Keywords:**

Artificial intelligence, automated news, local journalism, news literacy, social cohesion, empathy, public sphere, deep learning.

## **Introduction: scenarios of dynamic news narrative**

Automated news writing is becoming an important feature of news production, for example in sports, election results and real estate news. A plethora of solutions have been developed for highly structured data, that is, where all relevant information is stored and can be queried from a traditional database, and at the same time the goals of the system are precisely defined (Leppänen et al. 2017). The phenomenon has been increasingly analysed by researchers. Diakopoulos (2019, p. 4) calls it “The Algorithmic Evolution of News Media.” Furthermore, Miroshnichenko (2018) suggests four areas where computational, specifically artificial intelligence (AI) perspectives on news production have become ubiquitous: data mining in big data, selection of newsworthy topics, moderation of commentaries, and text writing. Beyond that, we have the concept of augmented journalism that includes AI tools to support the journalist's creativity (Lindén 2018, Maiden et al. 2020); identify suitable angles to a story (Opdahl and Tessem 2020); recommend news articles (Karimi et al. 2018) and personalise the reader experience (Svendsen et al. 2019).

Not all data have the highly structured and easily accessible forms used in present template-based systems. Take social media, for instance, where different sources of text, video, images and metadata permeate the virtual space. Advances in database technology accommodate such data formats as we depart from relational (tabular) databases into to non-relational (non-tabular) ones. Taking these developments into account, this chapter addresses news story production that uses not only highly structured information. What we call "dynamic news automation" will in principle be able to provide news stories from ongoing, unexpected, and surprising events. They are AI solutions that will select and organize story elements like photo, video, and text in real-time. We envision the use of deep learning, a class of machine learning approaches that uses deep neural network models (LeCun et al. 2015). Our chapter focuses on news tips, comments, and engagement from citizens, and how AI can support these features of journalism. We believe these activities are crucial to the quality of the public sphere in general and to local journalism in particular.

The research question for our chapter is: How can artificial intelligence (AI) support citizen engagement in news production for the local community? The chapter consists of a theoretical section describing the hypothetical design idea and the scenario approach adopted as our research method here. There is a description of an existing software ("Tipps") and three scenarios describing a progressively more advanced system for citizen contributions to local journalism. We foresee AI tools that may become particularly valuable for news automation

and its projected capabilities approximately five to ten years into the future. The scenarios describe how we may use existing knowledge for example about image recognition (LeCun et al. 2015), natural language processing (Young et al. 2018) and narrative construction (Gervas et al. 2019).

The scenarios are built around personas who contribute with news tips related to local transportation. Some people commute by car and are concerned with road safety and new road construction, others ride a bike to work. They are likely to have different interpretations of news relating to transportation. AI technology can select and bring forth opinions from multiple viewpoints, potentially contributing to mutual understanding and even empathy among contributors and news readers (Bazalgette 2017). The scenarios are intended to show that by utilizing AI in a news setting there can be better ways of engaging citizens in local journalism and the democratic public sphere than what current social media achieves.

## **Theory: Designing better publics**

The normative approach adopted here can be called hypothetical media design (Nyre 2014). Instead of analysing what has already happened, media design studies future possibilities by synthesizing and exploring communicative features that may become important in the future (Nyre 2014). This is in line with methods from other researchers, who have explored the possibility of transforming the public with new technological designs (Disalvo et al. 2014). It is a matter of identifying problems and challenges in the public and proposing actions that may solve them.

The design challenge we formulate is related to supporting citizen contributions to established local newspapers. These have an independent gatekeeping and agenda-setting role, and they often dominate the public sphere in their community. But local newspapers have had a difficult time staying relevant with the shift from reading the news on paper/browser to reading the news in social media. Some ask whether journalism has a future or not (Lepore 2019). Building advanced AI supported tools for automated real time journalism can be an answer to the problem.

Citizens have an important role in traditional journalism. Jürgen Habermas argued that citizens should keep their eyes and ears open and tip the public about problems relevant to the political system (Habermas 1996). The democratic ideal formulated for the public sphere has better technical preconditions than ever before. A range of interactive communication

technologies have allowed citizens to change from being passive audiences to becoming active contributors of journalistic content.

However, trolling, hate speech, divisive rhetoric, and fake news have become influential in the same historical process. There is concern about an increasing lack of mutual understanding among contributors to the public sphere suggested by terms like "filter bubbles" and "confirmation bias". All of this may lead to polarization or fragmentation, in which groups who may live in the same town or city have little common ground and few shared sources of information (Spinney 2017).

The suggested solution relies on a cultivation of empathy and perspective taking. Baron-Cohen (2004) defines empathy as "the drive to identify another person's emotions and thoughts and respond to them with an appropriate emotion." But the capacity for empathy, in general and toward any specific group, is not fixed. Storytelling is itself a good way of exercising and increasing empathy, in which the reader can relate to characters in the fictional or non-fictional story and be concerned about what happens to them. In fact, many forms of participation in the arts, including theatre going and musical appreciation, have been found to increase empathy and empathetic actions (Bazalgette 2017). Two important techniques that help to increase empathy, through storytelling, artistic practice, or other means, are *perspective taking* and *perspective giving*. Perspective taking is concerned with being able to view a situation from another's viewpoint. Even perspective giving (explaining your own viewpoint to others) can help to increase empathy, especially when those who communicate their perspective belong to groups who feel they are not heard.

The design challenge we formulate rests on the presumption that relevant citizen-provided content can make local journalism more supportive of the local democratic public and its role in the larger political system. The proposal is to transform and improve the existing mobile app "*Tipp*s" so that it encourages citizens to adopt an increasingly sophisticated communicative behaviour towards the local newspaper and the wider democratic public.

The design challenge consists of solving three problems related to atomized citizen contributions to local journalism: 1) News tips from readers are often of low technical quality with inaccurate information, no photos, and no news angle. Furthermore, they are provided asynchronously through e-mails or informally via the phone, 2) There is little or no peer-to-peer dialogue about local news among the readers in the newspapers. The newspapers' public sphere tends to be one-way, and discussions often take place in social media. 3) There is a lack of news literacy (Tully et al. 2021) among the readers, relating to the requirements of a

good tip as well as the ethics of a good public discussion. This makes people prone to fake news and trolling.

## **Normative scenario-building with AI**

Methodologically, we rely on a version of *scenario-building*. Durance and Godet (2010) state that a scenario is a means to represent future reality that masters the constraints of the present. They identify an approach called *normative scenarios* that fits with our purposes. "Normative scenarios are constructed from alternative images of the future which may be both desirable and feared, and are conceived in a retro-projective way". (Durance and Godet 2010, p.1489). Our design approach is normative in the sense that the scenarios are meant to support and cultivate citizen behaviour that is valuable for the democratic public using automation and other technical solutions based on AI.

The three scenarios are rooted in the existing news tips app “Tipps” (see <https://tipps.no>) that was launched for two Norwegian local newspapers in the autumn of 2020. The connection to existing public spheres and markets makes the scenarios both realistic and transparent. Some of the authors participated in the development of Tipps.

The app involves readers in the news cycle by encouraging them to take photos of unforeseen newsworthy events in their surroundings, and to curate a news story on the basis of the photo. The story can be further researched or published directly depending on its quality and relevance. Through Tipps, newspapers can present local newsworthy events almost immediately with the contribution of engaged citizens in their community.

A number of AI technologies can support an extended Tipps app in the future. Solutions based on recent developments in neural networks, or deep learning, will be the leveraging technologies for our scenarios. Deep learning is already widely applied to image recognition, but advances in deep learning model architecture and pre-training have made it possible to parallelize natural language processing and apply deep learning language models to different types of tasks (Wolf et al. 2020). For instance, instead of working through a sentence from beginning to end, a new method called *transformer* does not require this machine learning process to be sequential, and machine learning training times can be reduced through parallelization (Vaswani et al. 2017). These advances include a pre-training technique called BERT (Bidirectional Encoder Representations from Transformers), developed by Google, which improves the understanding of user searches (Devlin et al.

2018), but also other aspects of text interpretation, like sentiment analysis, entity recognition, co-referencing, and text coherence.

The three scenarios below require sensitive communication between citizens and technology. The advances in deep learning mentioned above, may make it possible for them to come true. Six technological abilities are particularly important - all of which may be provided by extensive use of deep learning:

- Identification of objects and persons: the ability to recognize the content of an image.
- Interpreting the situation: the ability to produce an interpretation of the event from identified items, texts, and facts;
- Generating content: the ability to construct new text about the situation;
- Understanding the deeper context of the situation: the ability to collect and combine other relevant content to create elaborate coherent narrative;
- Checking reliability: the ability to estimate the reliability of the information on the basis of photographs, text, and metadata;
- Checking ethical aspects of the story: the ability to ensure privacy and other ethical concerns during live publication

The scenarios develop our ambitious ideas: from a minimum viable product (MVP) called Tipps to an intelligent application which allows users to interact with the algorithms to produce news, interact with fellow citizens, and learn how to improve their contributions over time.

## **Scenario 1: Higher quality news tips with automation**

This design scenario is intended to solve problem 1: News tips from readers are often of low technical quality with inaccurate information, no photos, and no news angle. The overall normative goal for this scenario is for the tool to help citizens to make better news stories. Helen is a user of the system. One day when she walks her dog, a cyclist coming from a bike lane crossed a trafficked road, was hit by a car, and fell off the bicycle. Helen photographs the cyclist sitting on the ground, obviously in pain, and holding his left hand around his right wrist. A car is parked on the pavement and there is a person standing by the cyclist talking on the phone. An AI agent generates a provisional news text from Helen's first description, automated object recognition, and GPS information that Helen can improve on or submit directly to the news provider. The journalist receiving Helen's photo immediately sees the news value of the image, and publishes it on the web with a manually moderated variation of



Helen's text. For Helen this may be sufficient, but the system could also help the author to check whether the events and situations of the story are described in a coherent way by requesting or suggesting improvements to the image caption. There is a need for technology that analyses the image, connects to the user's suggested input, and assists in modifying or rewriting the text. The app should be able to search for specific location-dependent information indicated in keywords entered by the user and include it in the suggested text. On the newsroom side tools can filter the information for the journalist, possibly flag the information as potential fake news, and further help the editorial team to decide whether to publish the information or not.

Image summarization is an active area of AI research (e.g., deep learning), and in its simplest versions needs object recognition, i.e., identifying the category of objects that are present in an image. Software for this is already available, for instance in services provided by Google's Vision service or Microsoft's Azure. However, just identifying objects is not very interesting from a journalistic perspective, except that recognized objects may be used as tags for future search. We therefore also require the AI tool to be able to recognize relations among objects in an image (Haldekar et al. 2017) and activities performed in an image (Girish et al. 2020). When activity recognition and relation recognition algorithms can perform at a higher level than today, the AI agent will be able to extract information useful for automatically writing a text that may accompany a photo.

The Tipps tool may at this stage have information in the shape of extracted information from the image, more or less rudimentary text suggestions from the tipper, as well as other information the tipper finds relevant. One may also have access to computational news angle templates which add further structure to the text (Motta et al. 2020).

Texts generated from template-based systems lack the flow and elegance that is found in a human-written text. Within computational creativity approaches have been developed for writing text using search-based algorithms. They are based on assumed facts and create an exposition (most often text based) of the content taking into consideration the predefined constraints (Gervás et al. 2009, Gervás et al. 2019). So far, even in such tools, the written text is still rather static, and lacks liveliness and creativity.

A recent and promising approach to storytelling is story generation from deep learning through so called generative adversarial network (GAN) approaches (Liu et al. 2020). Today it is possible to provide such a system with a set of terms and other constraints, and then the system will construct a natural language story similar in style to what is found in previous stories. GAN algorithms can, if trained properly, mimic the voice of the newspaper. For

example, by training the deep learning network with texts in Nynorsk from the local newspaper. It is pertinent to ask whether this would be a biased algorithm, for instance, if the newspaper leans towards the right or left in the political sphere.

## **Scenario 2: Recommendation technologies and dialogic narrative**

This design scenario is intended to solve problem 2: Newspapers contain little or no peer-to-peer dialogue about local news among their readers. The newspapers' public sphere tends to be one-way, and discussions often take place in social media unrelated to the newspaper.

The AI system could help the newspaper to make the story more appealing to people who hold different values. Opinions could be included not just to appeal to a narrow value group (e.g., ban all cars) but also to include more moderate value groups (e.g., separate lanes for bicycles). Thus, the AI system could help news readers to broaden their understanding by including recommendations of more diverse content, for example stories of the same or similar events that pursue alternative opinions. The AI system should also promote a more active form of reading to counteract the citizens' tendency to ignore stories that do not fit with their own point of view. Diversity as a principle in recommendation systems may facilitate this (Yadav et al. 2020).

The editorial part of Tipps analyses the content received from Helen and may decide that more input is needed. Based on matches with analogical events in the newspaper's database a plan for further follow up of the event is made. Existing content in the form of text (like weather reports), images, and video clips that may be reused for the development of this story is collected. A preliminary, dynamic narrative of the event is constructed, including references to similar events at the same place or in other places. The Tipps editor then asks Helen or other Tipps citizen journalists to contribute with photos, video, and textual notes about the event. For example, Tipps may deduce by analogy to a previous story that ambulance and ambulance personnel are to be expected at the scene and asks Helen to take a photo of the ambulance.

Roger appears at the scene and is informed by the Tipps AI that images from the site are welcome. Roger suddenly sees that the probable driver of the car holds his hands to his chest and sits down, obviously in pain. Roger takes a chaotic photo that includes the cyclist, car driver, ambulance personnel and the ambulance. He also adds a few keywords about what he thinks has happened, like "collision", and "heart attack". The news story is updated immediately.

Helen has been concerned about the crossing of cycle lane and road for a long time. Cyclists often pass at high speed, and there are no hindrances in place to reduce the speed of cyclists or cars. So, to express her concerns she takes a picture of the crossing and writes a short comment about how dangerous she thinks it is and argues that there is a need for municipal investments to reduce risks of events like this.

All this new content is continuously narrated into a description of the event on a web page, with short text explanations, images, video clips, and links to other web sites and previous stories. As soon as Helen complains about the crossing, the system provides links to related audience opinions, relevant documents from the local government, and articles about politicians who have expressed their positions on the matter.

All this may be supported by case-based planning approaches (Borrajo et al. 2015). As new content is added a narrative is built by finding similar narratives and modifying these. Extensions to deep learning of word embeddings (Mikolov et al. 2013) and document embeddings (Le and Mikolov 2014) may provide effective ways to find similar reports in the newspaper database. Embeddings in this context are many-dimensional vectors of real numbers that may represent a word, a document, or perhaps even an event. The dimensions of such vectors are in the hundreds or even thousands. The use of deep learning transformer models is researched to support coherent ordering of sentences (Cui et al. 2020). In the future we should expect to see such techniques applied to provide coherent and consistent narratives of multi-modal content about events.

### **Scenario 3: Higher journalistic literacy with detectors and nudgers**

This design scenario is intended to solve problem 3: There is a lack of journalistic literacy among the readers, relating to the requirements of a good tip as well as the ethics of a good public discussion. This makes people prone to fake news and trolling.

As Helen takes her picture and writes the text, Tipps analyses the photograph. Helen is advised to choose a different angle for the photo due to privacy issues. When Helen writes a comment to the story suggesting stricter traffic regulations at the site she is advised to comment critically and constructively on similar and other views on the same issue. These come from the newspaper's database and are found by using recommender systems that ensure not only similar content or content adapted to the Helen's preferences, but also

opposite opinions from what she holds. Technologies for opinion mining (Zhang et al. 2017) combined with diversity recommenders (Chaney et al. 2018) could suggest such content.

The professional journalist would always reflect on privacy and confidentiality before publishing a story. This form of reflection may be learned by the citizen if the AI, by using object and activity recognition, suggests privacy problems in the image; or by transformer techniques that discover harassment or privacy in text or opinions submitted.

If Helen chooses to write a sarcastic comment in response to one of the commentaries, the system detects its negativity. Automatic moderation systems using recent deep learning transformer technologies (Wolf et al. 2020) flag Helen's comment as potentially problematic, and she is advised not to use irony and sarcasm in commenting on others' viewpoints.

This third version of the system would have a process through which citizens can reflect on and improve their news literacy skills. Photographs that are better composed and preserves privacy in a responsible way can be suggested from deep learning image analysis. Deep learning embedding techniques may indicate similar events or reports that can give the tipper ideas about how to write a caption in a better way.

We envision increased public responsibility among citizens based on the increased journalistic literacy that comes out of their use of this third Tipps version. A powerful example of such "nudge" technology is Litterati (<https://www.litterati.org>). This is a tool which allows users to submit photographs of litter and allow communities to map how dirty their neighbourhoods are compared to others, and this can in turn make users become more environmentally minded.

## **Conclusion: AI tools will transform journalism**

We have addressed the research question: How can artificial intelligence (AI) support citizen engagement in news production for the local community? Using a scenario-building method we have come up with a consensus-oriented mobile app based on 1) news tip quality support, 2) an empathic dialogue between citizens and 3) literacy-enhancing procedures.

The journalistic tool that we envision could be within reach of artificial intelligence methods in ten years. It would be based on a number of AI approaches supporting computational creativity and computational narrative. They will be based on natural language understanding, natural language generation, image and video analysis, and flexible pattern-based narrative construction. Some of these techniques are already available and being commercialized in their current forms. Others still need improvement and adaptation to the

scenarios presented. But given the speed of technological developments in AI it is highly likely that it will be possible to develop the technologies we suggest.

The main barrier to our scenarios for innovation is probably to be found in the news organisations themselves. Some journalists may be worried that they contribute to creating a tool that will end their careers in the years to come, and it may take a while before newspapers are able to accept such a level of automation in content creation.

Anyhow, the automation of dynamic content production will be valuable for the skilled journalist of the future. Overall, we envision AI with deep learning technology to become an agent that transforms the interaction between journalists and the community. We envision it not as a disruptive technology that would end jobs and lead to shut-down of local newspapers, but as a transforming asset that drives revenue and readership and encourages local public debate.

**Acknowledgements.** This work was supported by the Norwegian Media Authority and the regional Norwegian innovation fund UH-nett Vest. It was also supported by industry partners and the Research Council of Norway through MediaFutures: Research Centre for Responsible Media Technology and Innovation, project number 309339.

## References

- Baron-Cohen, S. (2004). *The Essential Difference*. Penguin Adult.
- Bazalgette, P. (2017). *The Empathy Instinct: How to Create a More Civil Society*. Hachette UK.
- Borrajo, D., Roubíčková, A., & Serina, I. (2015). Progress in Case-Based Planning. *ACM Computing Surveys*, 47(2), 35:1–35:39. <https://doi.org/10.1145/2674024>
- Chaney, A. J. B., Stewart, B. M., & Engelhardt, B. E. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems* (pp. 224–232). New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3240323.3240370>
- Cui, B., Li, Y., & Zhang, Z. (2020). BERT-enhanced Relational Sentence Ordering Network. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language*

- Processing (EMNLP)* (pp. 6310–6320). Presented at the EMNLP 2020, Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.511>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]*. <http://arxiv.org/abs/1810.04805>. Accessed 23 June 2021
- DiSalvo, C., Lukens, J., Lodato, T., Jenkins, T., & Kim, T. (2014). Making public things: how HCI design can express matters of concern. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2397–2406). New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/2556288.2557359>
- Diakopoulos, N. (2019). *Automating the News*. Harvard University Press.
- Durance, P., & Godet, M. (2010). Scenario building: Uses and abuses. *Technological Forecasting and Social Change*, 77(9), 1488–1492. <https://doi.org/10.1016/j.techfore.2010.06.007>
- Gervas, P. (2009). Computational Approaches to Storytelling and Creativity. *AI Magazine*, 30(3), 49–49. <https://doi.org/10.1609/aimag.v30i3.2250>
- Gervás, P., Concepción, E., León, C., Méndez, G., & Delatorre, P. (2019). The long path to narrative generation. *IBM Journal of Research and Development*, 63(1), 8:1-8:10. <https://doi.org/10.1147/JRD.2019.2896157>
- Girish, D., Singh, V., & Ralescu, A. (2020). Understanding action recognition in still images. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (pp. 1523–1529). Presented at the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). <https://doi.org/10.1109/CVPRW50498.2020.00193>

- Habermas, J. (1996). *Between facts and norms: contributions to a discourse theory of law and democracy*. Cambridge, Mass: MIT Press.
- Haldekar, M., Ganesan, A., & Oates, T. (2017). Identifying spatial relations in images using convolutional neural networks. In *2017 International Joint Conference on Neural Networks (IJCNN)* (pp. 3593–3600). Presented at the 2017 International Joint Conference on Neural Networks (IJCNN). <https://doi.org/10.1109/IJCNN.2017.7966308>
- Karimi, M., Jannach, D., & Jugovac, M. (2018). News recommender systems - Survey and roads ahead. *Inf. Process. Manag.* <https://doi.org/10.1016/j.ipm.2018.04.008>
- Le, Q. V., & Mikolov, T. (2014). Distributed Representations of Sentences and Documents. *arXiv:1405.4053 [cs]*. <http://arxiv.org/abs/1405.4053>. Accessed 23 June 2021
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Lepore, J. Does Journalism Have a Future? *The New Yorker*. <https://www.newyorker.com/magazine/2019/01/28/does-journalism-have-a-future>
- Leppänen, L., Munezero, M., Granroth-Wilding, M., & Toivonen, H. (2017). Data-Driven News Generation for Automated Journalism. In *Proceedings of the 10th International Conference on Natural Language Generation* (pp. 188–197). Santiago de Compostela, Spain: Association for Computational Linguistics. <https://doi.org/10.18653/v1/W17-3528>
- Linden, C.-G. (2018). Algorithms Are a Reporter’s New Best Friend: News Automation and the Case for Augmented Journalism. *The Routledge Handbook of Developments in Digital Journalism Studies*, 237–250. <https://doi.org/10.4324/9781315270449>
- Liu, D., Li, J., Yu, M.-H., Huang, Z., Liu, G., Zhao, D., & Yan, R. (2020). A Character-Centric Neural Model for Automated Story Generation. *Proceedings of the AAAI*

*Conference on Artificial Intelligence*, 34(02), 1725–1732.

<https://doi.org/10.1609/aaai.v34i02.5536>

Maiden, N., Zachos, K., Brown, A., Apostolou, D., Holm, B., Nyre, L., et al. (2020). Digital creativity support for original journalism. *Communications of the ACM*, 63(8), 46–53.

<https://doi.org/10.1145/3386526>

Nyre, L. (2014). Media design method. *The Journal of Media Innovations*, 1(1), 86–109.

<https://doi.org/10.5617/jmi.v1i1.702>

Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. *arXiv:1310.4546 [cs, stat]*. <http://arxiv.org/abs/1310.4546>. Accessed 23 June 2021

Miroshnichenko, A. (2018). AI to Bypass Creativity. Will Robots Replace Journalists? (The answer Is “Yes”). *Information*, 9(7), 183. <https://doi.org/10.3390/info9070183>

Motta, E., Daga, E., Opdahl, A. L., & Tessem, B. (2020). Analysis and Design of Computational News Angles. *IEEE Access*, 8, 120613–120626.

<https://doi.org/10.1109/ACCESS.2020.3005513>

Opdahl, A. L., & Tessem, B. (2021). Ontologies for finding journalistic angles. *Software and Systems Modeling*, 20(1), 71–87. <https://doi.org/10.1007/s10270-020-00801-w>

Spinney, L. (2017). How Facebook, fake news and friends are warping your memory. *Nature*, 543(7644), 168–170. <https://doi.org/10.1038/543168a>

Svendsen, R. D., Gulla, J. A., & Frøland, J. (2019). Anbefaling av nyhetsinnhold i praksis. Fra algoritmer til personaliserte nyheter. (Recommending news content. From algorithms to personalised news). *Norsk Medietidsskrift*, 26(1), 1–22. In Norwegian.

Tully, M., Maksl, A., Ashley, S., Vraga, E. K., & Craft, S. (2021). Defining and conceptualizing news literacy. *Journalism*, 14648849211005888.

<https://doi.org/10.1177/14648849211005888>



- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention Is All You Need. *arXiv:1706.03762 [cs]*. <http://arxiv.org/abs/1706.03762>.
- Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent Trends in Deep Learning Based Natural Language Processing. *arXiv:1708.02709 [cs]*. <http://arxiv.org/abs/1708.02709>.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., et al. (2020). Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (pp. 38–45). Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-demos>.
- Yadav, N., Mundotiya, R. K., Singh, A. K., & Pal, S. (2021). Diversity in Recommendation System: A Cluster Based Approach. In A. Abraham, S. K. Shandilya, L. Garcia-Hernandez, & M. L. Varela (Eds.), *Hybrid Intelligent Systems* (pp. 113–122). Cham: Springer International Publishing. [https://doi.org/10.1007/978-3-030-49336-3\\_12](https://doi.org/10.1007/978-3-030-49336-3_12)
- Zhang, L., & Liu, B. (2017). Sentiment Analysis and Opinion Mining. In C. Sammut & G. I. Webb (Eds.), *Encyclopedia of Machine Learning and Data Mining* (pp. 1152–1161). Boston, MA: Springer US. [https://doi.org/10.1007/978-1-4899-7687-1\\_907](https://doi.org/10.1007/978-1-4899-7687-1_907)